ASYMPTOTIC BEHAVIOR OF RANDOM HEAPS

J. BEN HOUGH[1]

ABSTRACT. We consider a random walk W_n on the locally free group (or equivalently a signed random heap) with m generators subject to periodic boundary conditions. Let $\#T(W_n)$ denote the number of removable elements, which determines the heap's growth rate. We prove that $\lim_{n\to\infty} \frac{\mathbb{E}(\#T(W_n))}{m} \leq 0.32893$ for $m \geq 4$. This result disproves a conjecture (due to Vershik, Nechaev and Bikbov [4]) that the limit tends to $\frac{1}{3}$ as $m \to \infty$.

1. INTRODUCTION

We consider a random walk W_n on the locally free group $\mathrm{L}\mathcal{F}_m$ with m generators subject to periodic boundary conditions, $\langle g_0, g_1, \ldots, g_{m-1} : g_i g_j = g_j g_i \; \forall \; i \neq j \pm 1 \bmod m \rangle$. The random walk on this group has an elegant interpretation in terms of random heaps, which are defined in [3] and examined more carefully in [4]. Consider the lattice $\mathbb{Z}_m \times \mathbb{Z}^+$, where $\mathbb{Z}_m$ denotes the integers modulo m, and drop signed pieces with $+/-$ charges uniformly over the m columns. When a piece is dropped in a column it falls as far as possible subject to the condition that it cannot fall past a piece in its own column or a piece in either of the neighboring columns, i.e. the pieces have "sticky corners". Also, if a piece with a $+$ charge lands directly on top of a piece with a $-$ charge, the two pieces annihilate. More precisely, each piece in the random heap may be described by three coordinates: its height, horizontal position and sign. Let (h_j, j, s_j) denote the highest piece in column j, and set $h_j = 0$ if no such piece exists. When a new piece is to be added to the heap, its horizontal position k and sign s are chosen uniformly over the $2m$ possiblities. If columns $k-1$, k and $k+1$ are empty, then the height of the new piece is 1. If $h_k > \max\{h_{k-1}, h_{k+1}\}$ and $s_k \neq s$, the new piece and (h_k, k, s_k) will both annihilate. Otherwise, the new piece is added to the heap and its height is $\max\{h_{k-1}, h_k, h_{k+1}\} + 1$. A typical random heap is illustrated in Figure 1.
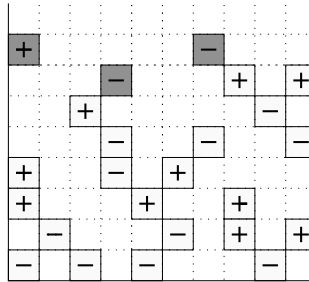

Figure 1. A signed random heap with 10 generators. The roof elements are shaded.

The interpretation of this construction should be clear. Dropping a $+$ (resp. $-$) piece in the k^{th} column corresponds to adding the generator g_k (resp. g_k^{-1}) to the random walk. The fact that g_k and g_{k+1} don't commute is reflected in the fact that pieces have sticky corners. In section 2, we describe a bijective

correspondence between signed heaps with m generators and reduced words in $\mathrm{L}\mathcal{F}_m$ so that a random walk on $\mathrm{L}\mathcal{F}_m$ corresponds to the growth procedure of the random heap detailed above. Our bijection differs from the one presented in [4], which is for the locally free group without periodic boundary conditions. This formalism allows us to use the terms random heap process and random walk on $\mathrm{L}\mathcal{F}_m$ interchangeably.

This geometrical picture of the random walk on $\mathrm{L}\mathcal{F}_m$ leads one to consider the *roof* $T(W_n)$ of the random heap, which is the collection of removable pieces. A piece is said to be *removable* if there are no pieces lying above it in its own column or the two adjacent columns. We shall say that a piece x *blocks* piece y if x must be annihilated before y can become part of the roof. The roof is an important feature of the random heap because its size controls the expected rate of growth of the heap. Indeed, if the roof has cardinality k, then the next piece to fall will annihilate a piece in the heap with probability $\frac{k}{2m}$, and will increase the size of the heap with probability $\frac{2m-k}{2m}$. The long term rate of growth of the heap is thus controlled by the long term expected size of the roof. To make this notion precise, we define the *drift* of the random heap process to be $\zeta = \lim_{n\to\infty}\frac{1}{n}\mathbb{E}(\#W_n)$ where $\#W_n$ denotes the number of pieces in the heap. It follows easily from the above observation that $\zeta = 1 - \lim_{n\to\infty}\frac{\mathbb{E}(\#T(W_n))}{m}$, a result given in [4].

We prove the following:

Theorem 1.1. *For $m \geq 4$, $\lim_{n\to\infty}\frac{1}{m}\mathbb{E}(\#T(W_n)) \leq 0.32893$.*

This result is surprising, since there are heuristic reasons described in [4] for suspecting that the limit should converge to $\frac{1}{3}$ as $m \to \infty$. In particular, one may consider a random walk on the semi-group $\mathrm{L}\mathcal{F}_m^+$ which has generators $g_0, g_1, \ldots, g_{m-1}$ satisfying the relations above, but does not include the inverses of the g_i's. This random walk $\tilde{W}_n$ has an interpretation as an unsigned random heap process. Specifically, we drop pieces uniformly over m columns as before, but now we forbid pieces from annihilating. Restricting attention to the roof of the unsigned heap, it is easy to see that this process is an irreducible Markov chain and we have the following:

Proposition 1.2. *(Proved in [4]) If $\tilde{W}_n$ is the unsigned random heap process with m generators, then $\lim_{n\to\infty}\frac{1}{m}\#T(\tilde{W}_n) = \frac{1}{3}$ a.s.*

The proof given in [4] may be simplified. If one defines $X_{n,i}$ to be the number of roof elements in the i^{th} column after the n^{th} particle is dropped, this process is a Markov chain on $\{0,1\}$ with transition probabilities $p_{1,0} = 1 - p_{1,1} = \frac{2}{m}$ and $p_{0,1} = 1 - p_{0,0} = \frac{1}{m}$. The stationary distribution for this 2-state chain is $(2/3, 1/3)$, so $\lim_{n\to\infty}\mathbb{E}\left(\#T(\tilde{W}_n)\right) = \lim_{n\to\infty}\sum_{i=1}^{m}\mathbb{E}(X_{n,i}) = \frac{m}{3}$. The claim now follows from the ergodic theorem. This elegant proof was discovered by Gábor Pete. From this nice result for unsigned heaps one might suspect that the same stationary behavior should be exhibited by the signed process, at least in the limit $m \to \infty$, a conjecture expressed in [4].

The paper is organized as follows. First we describe the bijection between elements of $\mathrm{L}\mathcal{F}_m$ and random heaps. Then we prove that for $m \geq 4$,

$$(1) \qquad\qquad \limsup_{n\to\infty}\frac{1}{mn}\sum_{k=1}^{n}\#T(W_k) \leq 0.32893 \; a.s.$$

Finally we shall deduce that if one allows heaps to be infinite, the random heap process has a unique stationary distribution ν and the finite dimensional distributions of W_n converge to the f.d.d.'s of ν as $n \to \infty$. From this fact, we deduce the same upper bound for the space average, proving the theorem. We conclude with some open problems.

2. The Bijection between Heaps and Words

We begin by giving a formal definition of a heap, and introducing some notation.

Definition 2.1. A *heap* H is a finite union of pieces (v_i, w_i, σ_i) with $v_i \in \mathbb{Z}^+$, $w_i \in \mathbb{Z}_m$ and $\sigma_i \in \{+, -\}$, which satisfy the following conditions:

1. If $(v_\alpha, w_\alpha, \sigma_\alpha) \in H$ and $v_\alpha > 1$, then there exists $(v_\beta, w_\beta, \sigma_\beta) \in H$ such that $v_\beta = v_\alpha - 1$ and $w_\beta \in \{w_\alpha - 1, w_\alpha, w_\alpha + 1\}$.
2. If $(v_\alpha, w_\alpha, \sigma_\alpha), (v_\beta, w_\beta, \sigma_\beta) \in H$ satisfy $v_\alpha = v_\beta + 1$ and $w_\alpha = w_\beta$, then $\sigma_\alpha = \sigma_\beta$.
3. If $(v_\alpha, w_\alpha, \sigma_\alpha), (v_\beta, w_\beta, \sigma_\beta) \in H$ satisfy $w_\alpha = w_\beta + 1$, then $v_\alpha \neq v_\beta$.

If H is a heap, the elements of the roof may all be removed to construct a new heap H'. We shall say that the roof of H' is the 2^{nd} level roof of H. Similarly, we may define the 3^{rd} level roof and so forth. It is easy to see that a heap may be determined uniquely by specifying the horizontal positions of the pieces in each of its roofs.

The concept of a roof also makes sense for words representing elements of $\mathrm{L}\mathcal{F}_m$. Specifically, if a is an instance of g_i or g_i^{-1} in the word W, we say that a is *removeable* if it commutes with all the letters in W occuring to its right, and no other instances of g_i or g_i^{-1} occur to its right. The *roof* of W consists of the collection of removeable letters. Omitting the removeable letters from W to obtain W', we define the roof of W' to be the 2^{nd} level roof of W. Higher level roofs are defined analogously. If $W = e$, the unit in $\mathrm{L}\mathcal{F}_m$, we say that all the roofs are empty. Observe that two words with the same roofs must represent the same element in $\mathrm{L}\mathcal{F}_m$, although words with different roofs may also be equal. To obtain a unique representation, we introduce the following definition.

Definition 2.2. A word W representing an element of $\mathrm{L}\mathcal{F}_m$ is said to be in *normal form* if the following conditions are satisfied. Here a is an instance of g_i or g_i^{-1} in W, and b is an instance of g_j or g_j^{-1}.

1. If $ab = e$, then a and b are not members of the same or adjacent roofs.
2. If a and b are members of the same roof and $i < j$, then b occurs to the right of a in W.
3. If a is a member of a higher level roof than b, then b occurs to the right of a in W.

The fact that every element of $\mathrm{L}\mathcal{F}_m$ may be represented by a unique normalized word follows directly from the commutation relations defining $\mathrm{L}\mathcal{F}_m$. By identifying the respective roofs of a heap and a normalized word so that a $+$ (resp. $-$) piece in column k corresponds to an instance of the generator g_k (resp. g_k^{-1}) one obtains the desired bijection, see Figure 2. With this identification, the random growth process of a heap described in the introduction corresponds to a random walk on $\mathrm{L}\mathcal{F}_m$.
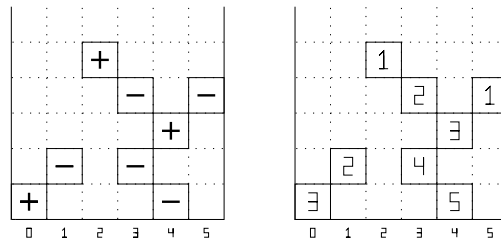

Figure 2. The pieces in the second figure are numbered to indicate their roof membership. The word corresponding to this heap is $g_4^{-1} g_3^{-1} g_0 g_4 g_1 g_3 g_2 g_5$.

3. Bounding the Time Average

We now give a proof of inequality (1). Consider a single column of a random heap, and when each piece is dropped record a "1" if the column contains an element of the roof and a "0" otherwise. This generates a sequence of 0's and 1's, say $X_1, X_2, X_3, \ldots$, and to prove (1) it is enough to show that $\limsup_{n \to \infty} \frac{1}{n} \sum_{k=1}^{n} X_k \leq$

0.32893 a.s. The sequence $X_1, X_2, X_3 \ldots$ may be expressed more compactly as $\langle 0 \rangle_0, \langle 1 \rangle_1, \langle 0 \rangle_1, \langle 1 \rangle_2, \langle 0 \rangle_2, \ldots$ where $\langle 1 \rangle_k$ denotes the length of the k^{th} sequence of 1's and similarly for $\langle 0 \rangle_k$. The random variable $\langle 0 \rangle_0$ is distinguished in that its value may be zero. Our approach is to show that

$$(2) \qquad \limsup_{n \to \infty} \frac{\sum_{k=1}^n \langle 1 \rangle_k}{\sum_{k=1}^n \langle 1 \rangle_k + \sum_{k=1}^{n-1} \langle 0 \rangle_k} \le 0.32893 \text{ a.s.}$$

We shall do this by constructing a.s. upper and lower bounds for

$$(3) \qquad \limsup_{n \to \infty} \frac{1}{n} \sum_{k=1}^\infty \langle 1 \rangle_k \quad \text{and} \quad \liminf_{n \to \infty} \frac{1}{n} \sum_{k=1}^n \langle 0 \rangle_k,$$

respectively.

An observation fundamental to our proof is that a sequence of 0's may be terminated in two different ways.

1. A piece may fall in the distinguished column.
2. All the pieces blocking the highest piece in the distinguished column may be annihilated.

We shall say that a 0 sequence *builds upward* if it is terminated via method 1 and otherwise it *backtracks*. Similarly, a sequence of 1's may be terminated in two ways.

1. A piece may fall in a column adjacent to the distinguished column.
2. All the pieces in the distinguished column blocking the highest piece in an adjacent column may be annihilated (or all the pieces in the distinguished column may be annihilated if the adjacent columns contain no pieces).

As before, we say that a sequence of 1's builds upward if it is terminated by method 1, and that it backtracks otherwise.

The first task is to construct an a.s. upper bound for $\limsup_{n \to \infty} \frac{1}{n} \sum_{k=1}^n \langle 1 \rangle_k$, the following outlines our approach. The key idea is based on the observation above. In general, if a sequence of 0's builds upward we expect the subsequent sequence of 1's to be significantly shorter than if the sequence of 0's backtracks. For in the first case the subsequent sequence of 1's may be easily terminated via either backtracking or building upward. Whereas in the latter case the sequence of 1's will likely begin at a time when there is a stack of pieces directly below the roof element in the distinguished column so it will be much more difficult to terminate the sequence of 1's via backtracking. We will describe this phenomenon precisely in Lemma 3.1, then will show in Lemma 3.3 that 0 sequences usually build upward. Henceforth, we shall say that a sequence of 1's starts from a *short position* if it is initiated at a time when there is no piece directly below the roof piece in the distinguished column. Otherwise we say that the 1 sequence starts from a *long position*. Also, we shall refer to the distinguished column as column 0 and label the columns to the right $1, 2, 3, \ldots$ modulo m. Unless otherwise stated, we assume without comment that $m \ge 4$.

Lemma 3.1. $\mathbb{E}(\langle 1 \rangle_k | \langle 1 \rangle_k$ *starts from a short position*$) = (\sqrt{2} - 1)m$ *and* $\mathbb{E}(\langle 1 \rangle_k) \le \frac{m}{2}$.

Proof. Consider the following random walk procedure on $\mathbb{Z}$. Start the walk from $\ell > 0$, and between consecutive steps of the walk wait i.i.d. times τ_i with distribution $\mathbb{P}(\tau_i > s) = \left(\frac{m-3}{m}\right)^s$. Also, immediately prior to taking each step flip a coin and with probability 2/3 stop the walk. If the walk is not stopped, then move 1 unit to the left or right with equal probability. Stop the walk when it reaches the origin, if it has not been stopped previously. It is easy to see that this process is equivalent to the process that determines the length of $\langle 1 \rangle_k$. The position of the random walk corresponds to the number of consecutive pieces at the top of column 0. Setting $\ell = 1$ we find:
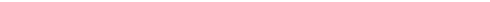
$$(4) \qquad \mathbb{E}(\langle 1 \rangle_k | \langle 1 \rangle_k \text{ starts from a short position}) = \mathbb{E}(\tau_i)\mathbb{E}(N \wedge T),$$

where $\mathbb{P}(N > k) = \left(\frac{1}{3}\right)^k$ and T is the hitting time of 0 for a simple random walk starting from 1. It follows from Lemma 3.3 in Chapter 3 of [2] that

$$(5) \qquad \mathbb{P}(T > 2k - 1) = \mathbb{P}(T > 2k) = \binom{2k}{k}\left(\frac{1}{4}\right)^k$$

so, since N and T are independent,

$$(6) \qquad \mathbb{E}(N \wedge T) = 1 + \sum_{k=1}^{\infty} \binom{2k}{k}\left(\frac{1}{4}\right)^k \left[\left(\frac{1}{3}\right)^{2k-1} + \left(\frac{1}{3}\right)^{2k}\right] = 3(\sqrt{2} - 1).$$

The above sum is evaluated using the identity $\sum_{k=0}^{\infty} \binom{2k}{k} t^k = (1 - 4t)^{-1/2}$ which is valid for $|t| < \frac{1}{4}$. Since $\mathbb{E}(\tau_i) = \frac{m}{3}$ we obtain $\mathbb{E}(\langle 1 \rangle_k) = (\sqrt{2} - 1)m$. From any starting position, $\mathbb{E}(\langle 1 \rangle_k) \leq \mathbb{E}(\tau_i)\mathbb{E}(N) = \frac{m}{2}$. $\qquad \square$

Figure 3. The shaded pieces in each figure are both realizations of the symbol ⬚⬚.

Next we bound the probability that a sequence of 0's backtracks. We say that a sequence of 0's starts from a short position if it is initiated at a time when there is only one piece blocking the highest piece in column zero. Then it is clear that the probability of backtracking is maximized when the 0 series starts from a short position. In what follows, we write

$$(7) \qquad \square\, ,\, \square\, ,\, \square\, ,\, \square\, ,\, \square\, ,\, \square\, ,\, \square\, ,\, \square\, ,\, \square\, ,\, \square\, ,\, \square\, ,\, \square\, ,\, \square$$

to denote possible configurations for the pieces blocking the uppermost piece in column zero. In each of the above pictures, the lowest piece is assumed to lie in column zero and we agree that exactly those pieces which block this piece are shown. The position of a piece indicates its column and which of the other pieces it blocks. It does not necessarily indicate its height relative to the others, see Figure 3. Two different realizations of a given symbol are said to represent the same *configuration* of blocking pieces. We shall use the notation $\mathbb{P}(⬚⬚)$ to denote the probability that a 0 sequence backtracks from the starting position ⬚⬚, and so forth. To see that this notation makes sense, we appeal to the following lemma.

Lemma 3.2. *The configuration of pieces blocking the highest piece in column zero determines both the probability of backtracking and the probability that this piece will be annihilated at any later time.*

Proof. Consider random heaps W^1 and W^2, and suppose that at time n the pieces blocking the highest piece in column zero of W_n^1 and W_n^2 have the same configuration. Label the highest pieces in column zero x_1 and x_2 respectively. Construct a coupling between W^1 and W^2 as follows. Select a piece uniformly over the $2m$ possible choices and add it to W_n^1 to obtain W_{n+1}^1. Now add a piece in the same column to W_n^2, and choose its sign according to the following rules.

1. If the newly added pieces both land on top of roof elements, fix the sign of the second piece so that both annihilate or neither annihilate.
2. Otherwise, choose the sign of the second piece to be the same as the sign of the first piece.

Observe that at time $n+1$, either both x_1 and x_2 will have been annihilated or the pieces blocking x_1 and x_2 will again have the same configuration. Iterating this procedure, the proof is complete. $\qquad\square$

We now estimate $\mathbb{P}(\square)$. Observe that by symmetry we have $\mathbb{P}(\square) = \mathbb{P}(\square)$.

Lemma 3.3. $0.137457 \leq \mathbb{P}(\square) \leq 0.14599$.

Proof. Conditioning on the next piece to fall in either column -1,0,1 or 2 we obtain:

$$(8) \qquad \mathbb{P}(\square) = \frac{1}{4}\mathbb{P}(\square) + \frac{1}{4}\cdot 0 + \frac{1}{8}\cdot 1 + \frac{1}{8}\mathbb{P}\left(\square\right) + \frac{1}{4}\mathbb{P}\left(\square\right).$$

Now expanding the terms on the RHS and using symmetry:

$$(9) \qquad \begin{aligned} \frac{291}{320}\mathbb{P}(\square) &= \frac{1}{8} + \frac{13}{160}\mathbb{P}\left(\square\right) + \frac{3}{20}\mathbb{P}\left(\square\right) + \frac{1}{20}\mathbb{P}\left(\square\right) + \frac{1}{20}\mathbb{P}\left(\square\right) \\ &\quad + \frac{1}{64}\mathbb{P}\left(\square\right) + \frac{1}{32}\mathbb{P}\left(\square\right) + \frac{1}{40}\mathbb{P}\left(\square\right). \end{aligned}$$

The lower bound for $\mathbb{P}(\square)$ is obtained by setting the unknown $\mathbb{P}(\cdot)$ terms on the RHS equal to zero. For the upper bound, we use the following lemma.

Lemma 3.4. *For any initial heap, call $\wp$ the probability that a specific piece in the roof is annihilated at any future time. Then $\frac{1}{6} < \wp < \frac{1}{5}$.*

Proof. Conditioning on the next piece to fall in any of the three relevant columns, we obtain

$$(10) \qquad \frac{1}{6} < \wp < \frac{1}{6} + \frac{5}{6}\wp^2.$$

Thus $(5\wp - 1)(\wp - 1) > 0$, and we deduce that $\wp < \frac{1}{5}$. $\qquad\square$

It follows that

$$\mathbb{P}\left(\square\right), \mathbb{P}\left(\square\right), \mathbb{P}\left(\square\right), \mathbb{P}\left(\square\right) \leq \frac{1}{25}\mathbb{P}(\square),$$

$$(11) \qquad \mathbb{P}\left(\square\right), \mathbb{P}\left(\square\right) \leq \frac{1}{5}\mathbb{P}(\square).$$

Substituting these approximations into (9) yields the upper bound. $\qquad\square$

The next task is to construct an a.s. upper bound on $\limsup_{n\to\infty} \frac{1}{n}\sum_{k=1}^{n} \langle 1 \rangle_k$ where $\langle 0 \rangle_0, \langle 1 \rangle_1, \langle 0 \rangle_1, \langle 1 \rangle_2, \ldots$ are chosen according to the heap process. We construct $\langle \tilde{1} \rangle_1, \langle \tilde{1} \rangle_2, \langle \tilde{1} \rangle_3, \ldots$ so that $\langle \tilde{1} \rangle_k \geq \langle 1 \rangle_k$ for all k and $\limsup_{n\to\infty} \frac{1}{n}\sum_{k=1}^{n} \langle \tilde{1} \rangle_k$ has a sufficiently small a.s. upper bound. The $\langle \tilde{1} \rangle_k$'s are constructed as follows. Consider the following four random variables: S, S^*, $\tilde{S}$ and L. The variable S shall be given the distribution of $\langle 1 \rangle_k$ starting from a short position. We give S^* the distribution of S conditioned on the event that the first piece to fall in either column -1, 0 or 1 after the 1 sequence is initiated actually falls in either column -1 or 1. Also, $\tilde{S}$ has the distribution of S conditioned on the event that the first piece to fall in either column -1, 0 or 1 after the 1 sequence is initiated actually falls in column 0. Finally, L has the distribution of a $\langle 1 \rangle_k$ whose starting position consists of an infinite tower of pieces in column 0. From Lemma 3.1, $\mathbb{E}(S) = (\sqrt{2} - 1)m$ and $\mathbb{E}(S^*) = \mathbb{E}(\tau_i) = \frac{m}{3}$, moreover $\mathbb{E}(L) = \mathbb{E}(N) = \frac{m}{2}$. Since $\frac{1}{3}\mathbb{E}(\tilde{S}) + \frac{2}{3}\mathbb{E}(S^*) = \mathbb{E}(S)$, we obtain $\mathbb{E}(\tilde{S}) = (3\sqrt{2} - 11/3)m$.

Suppose now that we have an infinite number of i.i.d. copies of S^*, $\tilde{S}$ and L. To construct $\langle \tilde{1} \rangle_k$, we first build a Markov chain ξ_k on the state space $\{s^*, \tilde{s}, \ell\}$, and then define $\langle \tilde{1} \rangle_k$ to be a new independent copy of S^*, $\tilde{S}$ or L if ξ_k equals s^*, $\tilde{s}$ or ℓ respectively. We shall ignore the condition $\langle \tilde{1} \rangle_k \geq \langle 1 \rangle_k$ for the time being. To

construct ξ_k, take $\xi_1 = s^*$ with probability 2/3 and $\xi_1 = \tilde{s}$ with probability 1/3. Once ξ_k has been chosen, choose ξ_{k+1} according to the following transition probabilities. Here $\rho = \mathbb{P}(\boxdot)$.

(12)
$$s^* \;\begin{array}{c} \nearrow \\ \rightarrow \\ \searrow \end{array}\; \begin{array}{ll} s^* & \frac{2}{3} \\ \tilde{s} & \frac{1}{3} \\ \ell & 0 \end{array} \qquad \tilde{s} \;\begin{array}{c} \nearrow \\ \rightarrow \\ \searrow \end{array}\; \begin{array}{ll} s^* & \frac{2}{3}(1-\rho) \\ \tilde{s} & \frac{1}{3}(1-\rho) \\ \ell & \rho \end{array} \qquad \ell \;\begin{array}{c} \nearrow \\ \rightarrow \\ \searrow \end{array}\; \begin{array}{ll} s^* & \frac{2}{3}(1-\rho) \\ \tilde{s} & \frac{1}{3}(1-\rho) \\ \ell & \rho \end{array}$$

Then ξ_k has stationary distribution:

(13)
$$\pi_{s^*} = \frac{2(1-\rho)}{3-2\rho}, \qquad \pi_{\tilde{s}} = \frac{1-\rho}{3-2\rho}, \qquad \pi_\ell = \frac{\rho}{3-2\rho}.$$

Since all the copies of S^* appearing among the $\langle \tilde{1} \rangle_k$'s are i.i.d., and similarly for the copies of $\tilde{S}$ and L, it follows from the ergodic theorem that

(14)
$$\frac{1}{n}\sum_{k=1}^{n} \langle \tilde{1} \rangle_k \xrightarrow{a.s.} \pi_{s^*}\mathbb{E}(S^*) + \pi_{\tilde{s}}\mathbb{E}(\tilde{S}) + \pi_\ell\mathbb{E}(L).$$

Combining equations (13) and (14) with Lemma 3.3 we obtain

(15)
$$\limsup_{n\to\infty} \frac{1}{n}\sum_{k=1}^{n} \langle \tilde{1} \rangle_k \le 0.41884m \text{ a.s.}$$

Now, to deduce the same a.s. bound for $\limsup_{n\to\infty} \frac{1}{n}\sum_{k=1}^{n} \langle 1 \rangle_k$, it remains only to demonstrate that the $\langle \tilde{1} \rangle_k$'s can be constructed so that $\langle \tilde{1} \rangle_k \ge \langle 1 \rangle_k$.

Lemma 3.5. *We may construct the sequences ξ_k and $\langle \tilde{1} \rangle_k$ with the joint densities specified above so that $\langle 1 \rangle_k \le \langle \tilde{1} \rangle_k$, and if $\xi_k = s^*$, then after $\langle 1 \rangle_k$ is initiated a piece falls in either column 1 or -1 before one falls in column 0.*

Proof. The proof is by induction. For the base case, take $\langle \tilde{1} \rangle_1 = \langle 1 \rangle_1$. To determine ξ_1, consider the first piece to fall in either columns -1, 0 or 1 after $\langle 1 \rangle_1$ is initiated. If it falls in either column -1 or 1, then $\xi_1 = s^*$, otherwise $\xi_1 = \tilde{s}$. Now assume that $\xi_1, \ldots, \xi_n$ and $\langle 1 \rangle_1, \ldots, \langle 1 \rangle_n$ have been constructed to satisfy the stipulations specified above. If $\xi_n = s^*$ we see by the induction hypothesis that $\langle 1 \rangle_{n+1}$ must start from a short position regardless of whether $\langle 0 \rangle_n$ backtracks. So take $\langle \tilde{1} \rangle_{n+1} = \langle 1 \rangle_{n+1}$ in this case. Define ξ_{n+1} as before by considering the first piece to fall in either column -1, 0 or 1 after $\langle 1 \rangle_{n+1}$ is initiated.

If $\xi_n = \tilde{s}$ or $\xi_n = \ell$, then construct a random variable χ_{n+1} independent of $\langle \tilde{1} \rangle_1, \ldots, \langle \tilde{1} \rangle_n$ and $\xi_1, \ldots, \xi_n$ so that $\mathbb{P}(\chi_{n+1} = 1) = 1 - \rho$ and $\mathbb{P}(\chi_{n+1} = 0) = \rho$. If $\mathcal{H}_n$ denotes the configuration of the heap at the time when $\langle 0 \rangle_n$ is initiated, we know from the discussion immediately preceding Lemma 3.3 that $\mathbb{P}(\langle 0 \rangle_n \text{ backtracks} \mid \mathcal{H}_n) \le \rho$. Thus, we may construct χ_{n+1} so that $\{\langle 0 \rangle_n \text{ backtracks}\} \subset \{\chi_{n+1} = 0\}$. Now, if $\chi_{n+1} = 1$, take $\langle \tilde{1} \rangle_{n+1} = \langle 1 \rangle_{n+1}$. Construct ξ_{n+1} as before by considering the first piece to fall in either column -1, 0 or 1 after $\langle 1 \rangle_{n+1}$ is initiated. If $\xi_n = 0$, take $\xi_{n+1} = \ell$. To construct $\langle \tilde{1} \rangle_{n+1}$, start building a random heap process on top of an infinite stack of pieces in the 0 column. Drop pieces in the same order as they fell while $\langle 1 \rangle_{n+1}$ was being constructed until $\langle 1 \rangle_{n+1}$ was terminated. At this point, if the 0 column still contains a piece in the roof, continue to drop signed pieces uniformly over the m columns, but now independently of the random heap process. Take $\langle \tilde{1} \rangle_{n+1}$ to be the length of the sequence of 1's that is generated. It is clear that if $\langle 1 \rangle_{n+1}$ builds upward, it will equal $\langle \tilde{1} \rangle_{n+1}$, and otherwise it will be strictly smaller. $\square$

Thus, we have established the following estimate:

Lemma 3.6. $\limsup_{n\to\infty} \frac{1}{n}\sum_{k=1}^{n} \langle 1 \rangle_k \le 0.41884m$ *a.s.*

We now calculate an a.s. lower bound for $\liminf_{n\to\infty} \frac{1}{n}\sum_{k=1}^{n} \langle 0 \rangle_k$. The first step is to estimate the conditional expectation $\mathbb{E}(\langle 0 \rangle_k | \langle 0 \rangle_k$ starts from a short position). In what follows, we shall write $\mathbb{E}(\text{⊔⊓})$ to denote $\mathbb{E}(\langle 0 \rangle_k | \langle 0 \rangle_k$ starts from ⊔⊓).

Lemma 3.7. $0.85453m \le \mathbb{E}(\text{⊏⊐}) \le 0.86255m$.

Proof. Conditioning on the first piece to fall in either column -1,0,1 or 2 we obtain

$$(16) \qquad \mathbb{E}(\text{⊏⊐}) = \frac{m}{4} + \frac{1}{4}\mathbb{E}(\text{⊔⊓}) + \frac{1}{8}\mathbb{E}\left(\text{⊟}\right) + \frac{1}{4}\mathbb{E}\left(\text{⊏⊐}\right),$$

and expanding the terms on the RHS again yields:

$$(17) \qquad \frac{291}{320}\mathbb{E}(\text{⊏⊐}) = \frac{61m}{160} + \frac{1}{64}\mathbb{E}\left(\text{⊟}\right) + \frac{13}{160}\mathbb{E}\left(\text{⊔⊟}\right) + \frac{1}{32}\mathbb{E}\left(\text{⊏⊐}\right) + \frac{3}{20}\mathbb{E}\left(\text{⊔⊏⊐}\right)$$
$$+ \frac{1}{20}\mathbb{E}\left(\text{⊏⊐}\right) + \frac{1}{20}\mathbb{E}\left(\text{⊏⊐}\right) + \frac{1}{40}\mathbb{E}\left(\text{⊏⊐}\right).$$

Bounding the unknown terms on the RHS by m yields the upper bound.

For the lower bound, it is useful to first estimate the related quantities $\tilde{\mathbb{P}}(\text{⊏⊠})$ and $\tilde{\mathbb{E}}(\text{⊏⊠})$. Here, $\tilde{\mathbb{P}}(\text{⊏⊠})$ is the probability that starting from the indicated configuration, the piece distinguished by the $\times$ symbol will be annihilated before a piece falls in column 0. Similarly, $\tilde{\mathbb{E}}(\text{⊏⊠})$ gives the expected number of steps before either the distinguished piece is annihilated or a piece falls in column 0.

Lemma 3.8. *We have the estimates* $0.180115 \le \tilde{\mathbb{P}}(\text{⊏⊠}) \le 0.1806355$ *and* $0.133939 \le \tilde{\mathbb{P}}\left(\text{⊏⊠}\right) \le 0.141677$. *For* $m = 4$, $\tilde{\mathbb{P}}\left(\text{⊏⊠}\right) = \tilde{\mathbb{P}}(\text{⊏⊠})$ *and for* $m \ge 5$ *we have* $0.133939 \le \tilde{\mathbb{P}}\left(\text{⊏⊠}\right) \le 0.141677$.

Proof. Conditioning on the next piece to fall in column -1, 0, 1 or 2 we write:

$$\tilde{\mathbb{P}}(\text{⊏⊠}) = \frac{1}{4}\tilde{\mathbb{P}}(\text{⊔⊠}) + \frac{1}{4}\cdot 0 + \frac{1}{8}\cdot 1 + \frac{1}{8}\tilde{\mathbb{P}}\left(\text{⊟⊠}\right) + \frac{1}{4}\tilde{\mathbb{P}}\left(\text{⊏⊠}\right)$$
$$(18) \qquad = \frac{1}{4}\tilde{\mathbb{P}}(\text{⊏⊠}) + \frac{1}{8} + \frac{1}{8}\tilde{\mathbb{P}}(\text{⊏⊠})^2 + \frac{1}{4}\tilde{\mathbb{P}}\left(\text{⊏⊠}\right)\tilde{\mathbb{P}}(\text{⊏⊠}).$$

Expanding $\tilde{\mathbb{P}}\left(\text{⊏⊠}\right)$ by conditioning on the next piece to fall in column 0,1,2 or 3 gives:

$$(19) \qquad \tilde{\mathbb{P}}\left(\text{⊏⊠}\right) = \frac{1}{4}\tilde{\mathbb{P}}\left(\text{⊔⊠}\right) + \frac{1}{8} + \frac{1}{8}\tilde{\mathbb{P}}\left(\text{⊏⊠}\right)^2 + \frac{1}{4}\tilde{\mathbb{P}}\left(\text{⊏⊠⊐}\right).$$

The probability that the highest piece in ⊏⊠ is annihilated at all is $\wp \le 1/5$, and the probability that it is annihilated by the first piece to fall in columns 0, 1, 2, 3 or 4 is $1/10$ (actually with $1/10$ replaced by $1/8$ if $m = 4$, but this extra precision is not important). Thus, we obtain the inequality:

$$(20) \qquad \frac{1}{10}\tilde{\mathbb{P}}\left(\text{⊏⊠}\right) \le \tilde{\mathbb{P}}\left(\text{⊏⊠⊐}\right) \le \frac{1}{5}\tilde{\mathbb{P}}\left(\text{⊏⊠}\right).$$

The quantity $\mathbb{P}\left(\text{⊏⊠}\right)$ satisfies the same inequality. Combining these inequalities with (19) and (20), the bounds on $\tilde{\mathbb{P}}\left(\text{⊏⊠}\right)$ follow. Once these inequalities are established, the bounds on $\tilde{\mathbb{P}}(\text{⊏⊠})$ follow from (18). To bound $\tilde{\mathbb{P}}\left(\text{⊏⊠}\right)$ if $m \ge 5$, condition on the next piece to fall in column 0, 2, 3 or 4. Then use the analogues of (20) to bound the unknown higher order terms. $\square$

The quantity $\tilde{\mathbb{E}}(\square^{\boxtimes})$ may be readily calculated from $\tilde{\mathbb{P}}(\square^{\boxtimes})$. Indeed, let T be the number of steps until either the distinguished piece is annihilated or a piece falls in column zero. By applying the strong Markov property at T, we obtain

$$(21) \qquad \tilde{\mathbb{E}}\left(\square^{\boxtimes}\right) = \left[1 + \tilde{\mathbb{P}}(\square^{\boxtimes})\right] \tilde{\mathbb{E}}(\square^{\boxtimes}).$$

Iterating this procedure yields

$$(22) \qquad \tilde{\mathbb{E}}\left(\left.\begin{matrix}\square\\\vdots\\\square^{\boxtimes}\end{matrix}\right\}^{n}\right) = \left[1 + \tilde{\mathbb{P}}(\square^{\boxtimes}) + \cdots + \tilde{\mathbb{P}}(\square^{\boxtimes})^{n}\right] \tilde{\mathbb{E}}(\square^{\boxtimes}),$$

and letting $n \to \infty$ we deduce that

$$(23) \qquad \tilde{\mathbb{E}}(\square^{\boxtimes}) = \left(1 - \tilde{\mathbb{P}}(\square^{\boxtimes})\right) m.$$

Analogous expressions clearly hold for $\tilde{\mathbb{E}}\left(\square^{\boxtimes}\right)$ and $\tilde{\mathbb{E}}\left(\square^{\boxtimes}\right)$. Using (23) and conditionings similar to the one leading to (21), we are now able to give lower bounds for the unknown terms in (17).

$$(24) \qquad \mathbb{E}\left(\square\right) \;\geq\; \tilde{\mathbb{E}}\left(\square\right) = \left[1 + \tilde{\mathbb{P}}(\square^{\boxtimes}) + \tilde{\mathbb{P}}(\square^{\boxtimes})^{2}\right] \tilde{\mathbb{E}}(\square^{\boxtimes}) \geq 0.994106m$$

$$(25) \qquad \mathbb{E}(\square) \;\geq\; \tilde{\mathbb{E}}\left(\square\right) \geq \left[1 + \tilde{\mathbb{P}}(\square^{\boxtimes})\right] \tilde{\mathbb{E}}(\square^{\boxtimes}) \geq 0.96737m$$

$$(26) \qquad \mathbb{E}(\square) \;\geq\; \tilde{\mathbb{E}}\left(\square\right) \geq \tilde{\mathbb{E}}(\square^{\boxtimes}) + \tilde{\mathbb{P}}\left(\square^{\boxtimes}\right)\left[1 + \tilde{\mathbb{P}}(\square^{\boxtimes})\right] \tilde{\mathbb{E}}(\square^{\boxtimes}) \geq 0.995377m$$

$$(27) \qquad \mathbb{E}(\square) \;\geq\; \tilde{\mathbb{E}}\left(\square\right) \geq \tilde{\mathbb{E}}\left(\square^{\boxtimes}\right) + \tilde{\mathbb{P}}\left(\square^{\boxtimes}\right) \tilde{\mathbb{E}}(\square^{\boxtimes}) \geq 0.974408m$$

$$(28) \qquad \mathbb{E}(\square) \;\geq\; \tilde{\mathbb{E}}\left(\square\right) \geq \tilde{\mathbb{E}}(\square^{\boxtimes}) + \tilde{\mathbb{P}}(\square^{\boxtimes})\left[\tilde{\mathbb{E}}\left(\square^{\boxtimes}\right) + \tilde{\mathbb{P}}\left(\square^{\boxtimes}\right) \tilde{\mathbb{E}}(\square^{\boxtimes})\right] \geq 0.995377m$$

$$(29) \qquad \mathbb{E}\left(\square\right) \;\geq\; \tilde{\mathbb{E}}\left(\square^{\boxtimes}\right) \geq \tilde{\mathbb{E}}(\square^{\boxtimes}) + \tilde{\mathbb{P}}(\square^{\boxtimes})\left[\tilde{\mathbb{E}}\left(\square^{\boxtimes}\right) + \tilde{\mathbb{P}}\left(\square^{\boxtimes}\right) \tilde{\mathbb{E}}(\square^{\boxtimes})\right] \geq 0.995377m$$

$$(30) \qquad \mathbb{E}\left(\square\right) \;\geq\; \tilde{\mathbb{E}}\left(\square\right) \geq \tilde{\mathbb{E}}(\square^{\boxtimes}) + \tilde{\mathbb{P}}(\square^{\boxtimes})\left[\tilde{\mathbb{E}}\left(\square^{\boxtimes}\right) + \tilde{\mathbb{P}}\left(\square^{\boxtimes}\right) \tilde{\mathbb{E}}(\square^{\boxtimes})\right] \geq 0.996374m$$

The numerical estimates are generated by using the estimates from Lemma 3.8. Substituting these values into equation (17) now gives the lower bound for $\mathbb{E}(\square)$. $\qquad\square$

It remains to construct an a.s. lower bound for $\liminf_{n\to\infty} \frac{1}{n} \sum_{k=1}^{n} \langle 0 \rangle_k$. Let $\langle 0 \rangle$ be a random variable with distribution equal to $\langle 0 \rangle_k$ starting from a short position. Since the distribution of $\langle 0 \rangle_k$ given $\mathcal{H}_k$ is stochastically greater than $\langle 0 \rangle$, we can construct $\langle \tilde{0} \rangle_1, \langle \tilde{0} \rangle_2, \langle \tilde{0} \rangle_3, \ldots$ to be i.i.d. copies of $\langle 0 \rangle$ so that $\langle \tilde{0} \rangle_k \leq \langle 0 \rangle_k$ for all k. By Lemma 3.7 and the strong law of large numbers, we know that $\limsup_{n\to\infty} \frac{1}{n} \sum_{k=1}^{n} \langle \tilde{0} \rangle_k \geq 0.85453m$ a.s. Hence, we arrive at the estimate

Lemma 3.9. $\limsup_{n\to\infty} \frac{1}{n} \sum_{k=1}^{n} \langle 0 \rangle_k \geq 0.85453m$ *a.s.*

It follows that

$$(31) \qquad \limsup_{n\to\infty} \frac{1}{n} \sum_{k=1}^{n} X_k \;\leq\; \limsup_{n\to\infty} \frac{\sum_{k=1}^{n} \langle 1 \rangle_k}{\sum_{k=1}^{n} \langle 1 \rangle_k + \sum_{k=1}^{n-1} \langle 0 \rangle_k}$$

$$\leq \frac{0.41884}{0.41884 + 0.85459} \leq 0.32893.$$

By symmetry, this result is valid for any column, and equation (1) follows.

4. Bounding the Space Average

We have shown that $\limsup_{n\to\infty} \frac{1}{mn} \sum_{k=1}^{n} \#T(W_k) \leq 0.32893$ a.s. Since $0 \leq \#T(W_k) \leq m$ for all k, it follows that

$$(32) \qquad \limsup_{n\to\infty} \frac{1}{mn} \sum_{k=1}^{n} \mathbb{E}(\#T(W_k)) \leq 0.32893.$$

Thus, to prove that $\limsup_{n\to\infty} \frac{1}{m} \mathbb{E}(\#T(W_n)) \leq 0.32893$ it suffices to show that $\lim_{n\to\infty} \mathbb{E}(\#T(W_n))$ exists. We shall prove the much stronger result that if one allows heaps to be infinite, the random heap process has a unique stationary distribution ν and the finite dimensional distributions of W_n converge to the corresponding f.d.d.'s of ν in total variation distance as $n \to \infty$.

We begin by introducing some notation. We shall write $T_k(W)$ to denote the configuration of the roofs of W above the $(k+1)^{st}$ level. Thus, $T_1(W) = T(W)$, $T_2(W)$ consists of the 1^{st} and 2^{nd} level roofs, etc. Infinite heaps are determined abstractly by specifying the configuration of their roof at each level. The roofs of an infinite heap W_∞ must satisfy the same geometrical relations as those of a finite heap: $T_k(W_\infty)$ must determine a finite heap satisfying Definition 2.1 for all k. Finite heaps are a subset of these generalized heaps distinguished by the property that only finitely many of their roofs are nonempty. Since the uppermost k roofs of a heap are sufficient to determine the uppermost $k-1$ roofs after a piece is added, the random heap process extends naturally to a Markov operator, $\mathcal{P}$, on the space Ω_H of generalized heaps. In what follows R_k shall denote the space of possible configurations of the uppermost k roofs, and $\mathcal{M}(R_k)$ shall denote the metric space of probability measures on R_k with total variation distance. We work with the sigma field $\mathcal{F}_H$ on Ω_H generated by all sets of the form $T_k^{-1}(\phi)$, where $\phi \in R_k$ and k is a positive integer. If μ is a measure on Ω_H, it induces a measure $\mu \circ T_k^{-1}$ on R_k. We claim that $\mathcal{P}$ has a unique stationary distribution on the measure space $(\Omega_H, \mathcal{F}_H)$, a fact which will follow readily from the following lemma.

Lemma 4.1. *Let W^1 and W^2 be signed heaps starting from random initial configurations W_0^1 and W_0^2 with laws μ_0^1 and μ_0^2. Then $\mu_n^1 = \mathcal{P}^n \mu_0^1$ and $\mu_n^2 = \mathcal{P}^n \mu_0^2$, the laws of W_n^1 and W_n^2 respectively, satisfy $\|\mu_n^1 \circ T_\ell^{-1} - \mu_n^2 \circ T_\ell^{-1}\| \to 0$ as $n \to \infty$ uniformly in μ_0^1 and μ_0^2. Here $\|\cdot\|$ denotes total variation distance.*

Proof. Construct a coupling between W^1 and W^2 as in Lemma 3.2. Now suppose that at some point pieces fall successively in columns $0, 1, 2, \ldots, m-1$ and that the pieces falling in column 0 did not annihilate. Once this sequence is observed, the roofs of W^1 and W^2 must agree. Call Φ_k the event that this sequence appears k consecutive times. In order for $T_\ell(W_n^1)$ and $T_\ell(W_n^2)$ to disagree after the event $\Phi_{k+\ell}$ is observed, at least mk of the pieces comprising the $\Phi_{k+\ell}$ sequence must be annihilated. By conditioning sequentially on these pieces and using Lemma 3.4, this probability is bounded above by $\left(\frac{1}{5}\right)^{mk}$. Now, given $\epsilon > 0$ we may choose N, K large enough so that $\left(\frac{1}{5}\right)^{Km} < \frac{\epsilon}{2}$ and the probability that $\Phi_{K+\ell}$ is not observed in the first N steps is less than $\frac{\epsilon}{2}$. Since $\|\mu_n^1 \circ T_\ell^{-1} - \mu_n^2 \circ T_\ell^{-1}\|$ is bounded by the probability that $T_\ell(W_n^1)$ and $T_\ell(W_n^2)$ differ it follows that for all $n \geq N$, $\|\mu_n^1 \circ T_\ell^{-1} - \mu_n^2 \circ T_\ell^{-1}\| < \epsilon$. $\qquad\square$

Taking $\mu_0^1 = \mu_0$ and $\mu_0^2 = \mu_j$ the lemma implies that as $n \to \infty$, $\|\mu_n \circ T_\ell^{-1} - \mu_{n+j} \circ T_\ell^{-1}\| \to 0$ uniformly in j. Thus $\mu_n \circ T_\ell^{-1}$ converges in $\mathcal{M}(R_\ell)$ to a measure ν_ℓ, defined by $\nu_\ell(A) = \lim_{n\to\infty} \mu_n \circ T_\ell^{-1}(A)$ for all $A \in R_\ell$, which satisfies $\nu_\ell \circ T_\ell \circ T_k^{-1} = \nu_k$ for $k < \ell$. It follows from Kolmogorov's extension theorem that there is a unique measure ν on Ω_H satisfying $\nu \circ T_k^{-1} = \nu_k$ for all k. This measure is the unique stationary measure for $\mathcal{P}$. To see that $\mathcal{P}\nu = \nu$, observe that $\mathcal{P}$ induces a continuous function $\tilde{\mathcal{P}} : \mathcal{M}(R_k) \to \mathcal{M}(R_{k-1})$. Since $\tilde{\mathcal{P}}(\mu_n \circ T_k^{-1}) = \mu_{n+1} \circ T_{k-1}^{-1}$ and $\|\mu_n \circ T_k^{-1} - \nu_k\|, \|\mu_n \circ T_{k-1}^{-1} - \nu_{k-1}\| \to 0$ as $n \to \infty$, continuity implies that $\tilde{\mathcal{P}}(\nu_k) = \nu_{k-1}$. Hence, $(\mathcal{P}\nu) \circ T_{k-1}^{-1} = \nu_{k-1}$ and it follows that $\mathcal{P}\nu = \nu$. Thus, ν is a stationary measure and by Lemma 4.1 it is unique.

5. Concluding Remarks and Further Problems

1. We have shown that for a signed random heap, $\frac{1}{m} \lim_{n \to \infty} \mathbb{E}(\#T(W_n)) \leq 0.32893$. However, the exact value of the limit is unknown for any $m \geq 4$.

2. The random heap process may be generalized as follows. Consider dropping pieces with sticky corners uniformly over the m columns, and when a piece lands directly on top of another piece, flip a coin and allow the two pieces to annihilate with probability p. When $p = 0$, this new process is equivalent to the unsigned random heap process, and if $p = \frac{1}{2}$, we recover the signed random heap process. Taking $p = 1$, we obtain a random walk on the group

$$\left\langle g_0, g_1, \ldots, g_{m-1} : g_i g_j = g_j g_i \ \forall i \neq j \pm 1 \bmod m, \text{ and } g_i^2 = e \ \forall i \right\rangle.$$

Numerical simulations for small values of m suggest that the expected size of the roof is monotonic in p, but proving this remains a challenge.

References

1. Desbois, J. and Nechaev, S. Statistics of reduced words in locally free and braid groups: abstract studies and applications to ballistic growth model, *J. Phys.* (A): Math. Gen. **31** (1998), 2767-89.

2. Durrett, R. *Probability: Theory and Examples*, 2nd ed, Duxbury Press, Belmont, 1996.

3. Vershik, A. M. Dynamic theory of growth in groups: entropy, boundaries, examples, *Russian Math. Surveys* **55**, No.4 (2000), 667-733.

4. Vershik, A. M., Nechaev, S. and Bikbov, R. Statistical properties of locally free groups with applications to braid groups and growth of random heaps, *Commun. Math. Phys.* **212** (2000), 469-501.